

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets

(11)



EP 0817 445 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.01.1998 Bulletin 1998/02

(51) Int Cl. 6: H04L 29/06, H04L 12/56

(21) Application number: 97304781.4

(22) Date of filing: 01.07.1997

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE

(30) Priority: 02.07.1996 US 678408

(71) Applicant: SUN MICROSYSTEMS, INC.
Mountain View, CA 94043 (US)

(72) Inventor: Katiyar, Dinesh
Mountain View, California 94040 (US)

(74) Representative:
Cross, Rupert Edward Blount et al
BOULT WADE TENNANT,
27 Furnival Street
London EC4A 1PQ (GB)

(54) Apparatus and method for identifying server computer aggregation topologies

(57) A method of processing a remote procedure call from a client computer to an object stored on an aggregation of server computers includes the step of checking a server aggregation location data field and a server aggregation contact strategy data field of the remote procedure call. The aggregation of server comput-

ers is designated as replicating server computers, migrating server computers, or federated server computers based upon the checking operation. Once a server computer aggregation topology is identified, parameters associated with the remote procedure call may be modified to alter the interaction with the server computer aggregation.

BEST AVAILABLE COPY

EP 0817 445 A2

Description**Brief Description of the Invention**

This invention relates generally to aggregations of server computers that collaborate to provide a service in distributed client-server computer systems. More particularly, this invention relates to a technique of interacting with an aggregation of server computers based upon a derived characterization of the server computer aggregation.

Background of the Invention

In a client/server computer network, the user of a client computer requests the execution of an object. In particular, the user requests the execution of a method associated with the object. Frequently, the object is not stored locally on the client computer. Thus, a remote procedure call (RPC) must be made to a server computer on which the object resides. In most cases, the server computer executes the requested method to generate a result. The result is then passed to the client computer.

It is common to provide an aggregation of server computers that collaborate to supply a service to client computers. This collaboration of server computers may be referred to as server aggregation. Reasons to use server aggregation include: robustness and fault-tolerance, load-balancing, data partitioning, allowing for system evolution (upgrades, etc.), priority and secure servicing, etc. The development of almost any mission-critical application makes use of some form of aggregation, and thus it is important that any distributed application develop a framework that provides support for it.

Figure 1 illustrates a client/server computer apparatus 20 in accordance with the prior art. The apparatus 20 includes a set of client computers 22A-22N, which are each linked to a transmission channel 23. The transmission channel 23 generically refers to any wire or wireless link between computers. The client computers 22A-22N use the transmission channel 23 to communicate with a set of server computers 24A-24N, forming a server aggregation 25.

Each client computer 22 has a standard computer configuration including a central processing unit (CPU) 30, connected to a memory 32, which stores a set of executable programs. The executable programs in this exemplary system include at least one client application program 34, client stubs 38, client subcontracts 40, and an operating system 42.

The client application program 34 is any application-level program, such as an application program that the user of a client computer 22 interacts with. The client stubs 38 receive procedure calls by the application program 34 requesting the execution of specified methods of specified objects. The purpose of the client stubs 38 is to access objects that are implemented in other address spaces, such as at the server computers 24A-

24N.

The client subcontract programs 40 and server subcontract programs 58 control the basic mechanisms of object invocation and argument passing. They control how object invocation is implemented, how object references are transmitted between address spaces, how object references are released, and similar object runtime operations. For example, when a client invokes an object of a given subcontract, the subcontract implements the object invocation by transmitting the request to the address space where the associated object is located, commonly a server computer 24 of the server aggregation 25.

The client subcontract programs 40 perform a marshal operation to transmit an object invocation (i.e., a remote procedure call) to another address space. A corresponding un-marshalling operation is performed by a server subcontract 58 on a server computer 24. The client subcontract programs 40 also perform unmarshal operations when receiving a reply (such as the results generated from a method call) from another computer, say the server computer 24. An operating system 42 underlies the operations of the client application programs 34, the client stubs 38, and the client subcontracts 40.

Each server computer 24 has a configuration analogous to that of each client computer 22. Each server 24 includes a CPU 50 and an associated memory 52. The memory 52 stores server application programs 54, server stubs 56, server subcontract programs 58, and an operating system 60. As indicated above, the server stubs 56 handle incoming method invocations on an object and call the specified method to perform the operation. As also indicated above, the server subcontracts 58 perform data marshalling and other operations to support the transport of method invocations and the resulting return messages between the server 24 and the client computers 22.

The operations of the apparatus of Figure 1 are more fully appreciated with reference to Figure 2.

Figure 2 illustrates the client computer 22A, the transmission channel 23, and the server computer 24A. As indicated above, a client application 34 invokes a specified method of an object in a different address space using a remote procedure call. The remote procedure call is passed by the client stubs 38 to the client subcontract 40, which packages the remote procedure call for transport on the transmission channel 23. The server subcontract 58 of the server 24A receives the information and un-packages it. The server subcontract 58 then passes the information to the server stubs 56. The server stubs 56 access the server application programs 54, which are the previously described object methods. More specifically, a specified server stub 56 makes a procedure call to execute a specified method of the invoked object. The execution of the method produces a set of results, herein called a reply, which is then passed back to the server subcontract 58 and the communication path is reversed, as indicated by the arrows of Fig-

ure 2.

Block 70 of Figure 2 illustrates the components associated with an Object Request Broker (ORB). An ORB is a distributed mechanism for handling remote procedure calls. The mechanism is distributed in the sense that the software associated with an ORB is on both the client computer 22 and the server computer 24.

The present invention is directed toward a distributed object client/server computer system of the type described in relation to Figures 1 and 2. More particularly, the invention is directed toward the operation of a server computer aggregation 25 in a distributed object client/server computer system. Existing server computer aggregations lack a common architectural basis that can be used to provide clean and elegant implementations of server aggregation features. The absence of a common architectural basis for server aggregation features prevents a reasonably uniform view for both system users and administrators. Thus, it would be highly desirable to provide a clean and elegant implementation of server aggregation features. Such an implementation would provide a reasonably uniform systematic approach for both system users and administrators.

Summary of the Invention

An embodiment of the invention includes a method of processing a remote procedure call from a client computer to an object stored on an aggregation of server computers. A server aggregation location data field and a server aggregation contact strategy data field of the remote procedure call are checked. The aggregation of server computers is designated as replicating server computers, migrating server computers, or federated server computers based upon the checking operation. Once a server computer aggregation topology is identified, parameters associated with the remote procedure call may be modified to alter the interaction with the server computer aggregation.

The invention's identification and use of server aggregation topology information provides a clean and elegant implementation for server aggregation features. In addition, the server aggregation topology information provides a reasonably uniform and systematic approach for both system users and administrators.

Brief Description of the Drawings

Examples of the invention will be described in conjunction with the accompanying drawings, in which:

FIGURE 1 illustrates a prior art client/server computer system with a server computer aggregation.

FIGURE 2 illustrates the handling of remote procedure calls in the system of Figure 1.

FIGURE 3 illustrates processing operations associated with a client computer accessing a server computer aggregation with a dual-role proxy, in accordance with an embodiment of the invention.

FIGURE 4 illustrates processing operations associated with a client computer accessing a server computer aggregation with a dual-role proxy, in accordance with an embodiment of the invention.

FIGURE 5 illustrates processing operations associated with a client computer accessing a server computer aggregation with a dual-role proxy, in accordance with an embodiment of the invention.

FIGURE 6 illustrates processing operations associated with a client computer accessing a server computer aggregation with a dual-role proxy, in accordance with an embodiment of the invention.

FIGURE 7 illustrates an aggregation object reference data structure utilized in accordance with an embodiment of the invention.

FIGURE 8 illustrates a server computer configured in accordance with an embodiment of the invention.

FIGURE 9 illustrates a client computer configured in accordance with an embodiment of the invention.

FIGURE 10 illustrates a server aggregation characterization technique utilized in accordance with an embodiment of the invention.

Like reference numerals refer to corresponding parts throughout the several views of the drawings.

Detailed Description of the Preferred Embodiment

In the system of Figure 1, a client computer, say client computer 22A, contacts the aggregation of server computers 25 by initially contacting a primary server, say server computer 24A. More particularly, a client computer, say client computer 22A, generates a remote procedure call to an object stored on a primary server, say server computer 24A. (In the following discussion it is understood that a reference to a client computer 22 contacting a server computer 24 contemplates a remote procedure call from the client computer 22 to the server computer 24.) After the client computer accesses the primary server, other server computers, say server computers 24B and 24C, in the aggregation 25 are then contacted by the client computer 22A.

The present embodiment departs from the approach described in reference to Figure 1. In particular, the embodiment utilizes a dual-role smart proxy server computer as a gateway to an aggregation of server computers. This operation can be appreciated with reference to Figure 3.

Figure 3 illustrates a client computer 22A and a server computer aggregation 25 including server computers 24A-24N. In accordance with the invention, the server computer aggregation 25 includes three types of server computers: a dual-role proxy server computer 24A, a primary server computer 24B, and a set of secondary server computers 24C-24N.

The client computer 22A initially contacts the dual-role proxy server computer 24A. Thereafter, the dual-role proxy server computer 24A contacts the remaining server computers 24B-24N of the aggregation 25. Thus,

In this embodiment of the invention, the dual-role proxy server computer 24A operates as a front-end server computer for the server aggregation 25.

The dual-role proxy server computer 24A performs this operation for client computers that are not clients of the aggregation 25. In other words, a client computer is involved, but the client computer is a non-client of the server aggregation 25, meaning it has substantially limited knowledge and access to the server aggregation 25.

Figure 4 illustrates a server aggregation 25 being accessed by an aggregation client server computer 22B. The aggregation client server computer 22B has substantial knowledge and access to the server aggregation. As a result, it can bypass the dual-role proxy server computer 24A and directly contact the primary server 24B and secondary servers 24C-24N. This approach is feasible when server aggregation information can be statically determined and embedded in an object reference.

Server aggregation information is not always static. Frequently, the location of servers is dynamically configurable and updated. In this case, it is necessary for a client computer to obtain this dynamic information.

Figure 5 illustrates a processing sequence that can accomplish the required update of dynamic information for a client computer. Figure 5 illustrates a client computer 22B directly accessing a primary server 24B, as shown with arrow 60. The primary server 24B then passes dynamic information back to the aggregation client computer 22B, as shown with arrow 62. With the updated dynamic information, the aggregation client computer 22B may then directly access the secondary servers 24C-24N, as shown with arrows 64 and 66. Note once again that the dual-role proxy computer 24A is bypassed.

Fault-tolerant behavior in a server aggregation 25 requires the aggregation client computer 22B to be able to communicate with the servers in the aggregation even if the primary server 24B is unavailable. In this case, the dual-role proxy server 24A may be exploited.

Figure 6 includes an arrow 60 representing the aggregation client computer 22B attempting to contact the primary server 24B. If the primary server 24B is inoperative, then the aggregation client computer 22B contacts the dual-role proxy server 24A, as shown with arrow 72. The dual-role proxy server 24A responds by passing updated aggregation information back to the aggregation client computer 22B, as shown with arrow 74. Thereafter, the aggregation client computer 22B can directly access the secondary servers 24C-24N, as shown with arrows 76 and 78.

It can now be appreciated that the dual-role proxy server computer 24A of the invention performs two roles. First, it operates as a server aggregation gateway for non-client server computers. Second, it operates as an information agent for client server computers.

The operation of an embodiment of the invention

has now been described. Attention presently turns to a discussion of different techniques that may be used to implement the invention.

Figure 7 illustrates an aggregation object reference data structure (AORDS) 80 that may be used in accordance with the invention. Each AORDS 80 includes location information for the smart proxy server of the server aggregation. In addition, each AORDS includes an object key, including a server aggregation location and a server aggregation contact strategy. The server aggregation location information includes a field for the location information for the primary server and a field for the location information for the secondary servers. Finally, each AORDS includes a client key, which stores information that identifies whether the object associated with the AORDS is a server aggregation client.

A non-client of a server aggregation will have missing fields in the AORDS data structure. In other words, a remote procedure call associated with a non-client of a server aggregation will have information regarding the location of the smart proxy server, will have client key information indicating a non-client status, but will not have, or will have incomplete, object key information. As indicated above, the object key information includes server aggregation location information and server aggregation contact strategy information. The dual-role proxy server is relied upon by the non-client to have the primary server location information and the secondary servers location information and to effectuate communications with those servers.

A client of a server aggregation will have complete fields in the AORDS data structure. That is, smart proxy server location information will be available, client key information (indicating status as a client) will be available, and object key information will be available, namely location information for the primary server and location information for the secondary servers, and a server aggregation contact strategy. While the information may be complete, it may have to be periodically updated, as described in reference to Figures 5 and 6.

Figure 7 illustrates that the AORDS data structure may be handled, along with other method call data, by client subcontracts 40 on the client computer side of the client/server system. As indicated in the background section, a client application program 34 generates a remote procedure call. Thereafter, client stubs 38 operate as an interface to a client sub-contract 40. The client sub-contract 40 is used to transport the remote procedure call over a transmission channel. Reverse processing is then performed at the server side. The general handling of object data structures in an object request broker of a distributed client/server system is known in the art.

Figure 8 illustrates a server computer 90 in accordance with an embodiment of the invention. The server computer 90 includes a CPU 92 connected to a memory 94, which stores a set of executable programs. Figure 8 illustrates server applications 96, including a client key

interpreter 98. The client key interpreter 98 interprets the client key field of the AORDS 80. In other words, the client key interpreter 98 determines whether a remote procedure call is associated with a server aggregation client.

A server aggregation gateway module 100 is also stored in the memory 94. The server aggregation gate module 100 includes the logic for accessing the primary servers and secondary servers of the server aggregation when a non-client remote procedure call is being processed.

Another application stored in the memory 94 of the server computer 90 is an aggregation update module 102. As its name implies, this executable program provides updated information on the server aggregation. This module is invoked, for example, during the processing shown in Figure 6.

Figure 8 also illustrates server stubs 104 stored in memory 94. Server subcontracts 106 are also stored in memory 94. The server subcontracts 106 handle the previously described AORDS 80. The memory 94 also includes an operating system 108.

The AORDS concept disclosed in accordance with an embodiment of the invention is highly beneficial because it supports the invention's concept of a dual-role proxy server computer. The concept finds further utility in its role of identifying different types of server aggregations. In other words, in accordance with another embodiment of the invention, the AORDS information is scrutinized to determine what type of server aggregation is being invoked by a remote procedure call. This information can then be used to change selected parameters associated with the remote procedure call.

This operation is more fully appreciated with reference to Figure 9. Figure 9 illustrates a client computer 110 operated in accordance with an embodiment of the invention. The client computer 110 includes a CPU 112 and a memory 114. The memory 114 stores a set of executable programs including client application programs 116. The client application programs 116 include an aggregation parameter check module 118. The aggregation parameter check module 118 checks the server aggregation location data field and the server aggregation contact strategy data field of the AORDS 80. Based upon this operation, the nature of the server aggregation 25 is determined.

Figure 10 illustrates the different types of server aggregation topologies that are recognized in accordance with the invention. If the server aggregation location is known to the client computer (and its associated remote procedure calls) and the server aggregation contact strategy is also known to the client computer (and its associated remote procedure calls), then a replicating server aggregation is said to exist. A replicating server aggregation is one in which multiple servers maintain synchronized and consistent information. Thus, multiple clients can independently access replica servers for the same information.

It is impractical to expect a client to know the location of a server aggregation, yet not know the strategy to contact the server aggregation, thus the lower left corner of the box of Figure 10 is crossed out.

5 Another server aggregation topology that is recognized in accordance with the invention is a migrating server aggregation. A migrating server aggregation exists when the client computer has a contact strategy (the server aggregation contact strategy of the AORDS 80), 10 but does not have the server aggregation location (the server aggregation location of the AORDS 80). Most common scenarios of migration have the client computer 22 contacting the primary server and being forwarded to the new location of the server. Thus, the client computer knows the strategy (i.e. following the forwarding pointer), but not the location of the new server.

15 The final server aggregation topology that is recognized in accordance with the invention is a federated server aggregation. For this topology, the server aggregation location is known to the server and the contact strategy is located with the server, as shown in Figure 10. Almost all cases of federation involve the primary server knowing both the locations of the other servers in the federation as well as the logic for interacting with them and building a response.

20 The aggregation designator 120 is an executable program that checks the server aggregation location field and the server aggregation contact strategy field of the AORDS 80 of the invention. Based upon the information it observes, it provides a designation to the server aggregation, pursuant to the framework shown in Figure 10.

25 Thus, the present invention uses the server aggregation location information and server aggregation contact strategy of the AORDS 80 to identify a server aggregation topology. These information fields represent two critical bits of control that determine the different kinds of interaction between a client and an aggregation of servers. In accordance with the embodiment one may 30 structure the implementation of server aggregation features around these two pieces of information, providing a select set of strategies that can be employed for interacting with the servers in the aggregation. This advantage also applies to administrative interfaces.

35 40 45 A remote procedure call processor 122 may be used to alter parameters associated with a remote procedure call once the server aggregation topology is identified. For example, if a replicating server aggregation is identified, the RPC processor 122 may be used to prefetch a list of replica servers that can respond to the same method call. In this case, if a client computer fails in its initial attempt to contact a server, the client computer immediately accesses the pre-fetched list of replica servers that can respond to the method invocation.

50 55 After a migrating server aggregation is identified, the RPC processor 122 may be used to set parameters that limit the number of migrations that can be performed

to satisfy a method call.

Another example of how the RPC processor 122 may respond to the aggregation designation is to set context flags associated with the method call. For instance, the context flags may be used to determine which requests are propagated to the federation and which are necessarily serviced by the primary server only.

Figure 9 also illustrates a set of client stubs 124 of the type previously described. Figure 9 also illustrates client subcontracts 126 that handle the disclosed AORDS 80. Finally, Figure 9 illustrates a operating system 128 associated with the client computer 110.

Claims

1. A client/server computer apparatus, comprising:

a transmission channel; 20
an aggregation of server computers connected to said transmission channel; and
a client computer connected to said transmission channel, said client computer executing
a first routine to check a server aggregation location data field and a server aggregation contact strategy data field of a remote procedure call generated by said client computer, and 30
a second routine to designate said aggregation of server computers as replicating server computers, migrating server computers, or federated server computers based upon the result of said first routine.

2. The apparatus of claim 1 wherein said first routine designates said aggregation of server computers as replicating server computers when said server aggregation location data field is complete and said server aggregation contact strategy data field is complete.

3. The apparatus of claim 1 wherein said first routine designates said aggregation of server computers as migrating server computers when said server aggregation location data field is incomplete and said server aggregation contact strategy data field is complete.

4. The apparatus of claim 1 wherein said first routine designates said aggregation of server computers as federated server computers when said server aggregation location data field is incomplete and said server aggregation contact strategy data field is incomplete.

5. The apparatus of claim 1 wherein said client com-

puter executes a third routine to modify a remote procedure call parameter after said aggregation of server computers is designated.

5 6. A computer readable memory storing executable instructions for execution by a client computer of a client/server computer system including an aggregation of server computers, comprising:

10 a plurality of instruction sets stored in said computer readable memory, including

15 a first instruction set to check a server aggregation location data field and a server aggregation contact strategy data field of a remote procedure call generated by said client computer; and

20 a second instruction set to designate said aggregation of server computers as replicating server computers, migrating server computers, or federated server computers based upon said check operation of said first instruction set.

25 7. The computer readable memory of claim 6 wherein said first instruction set includes instructions to designate said aggregation of server computers as replicating server computers when said server aggregation location data field is complete and said server aggregation contact strategy data field is complete.

30 8. The computer readable memory of claim 6 wherein said first instruction set includes instructions to designate said aggregation of server computers as migrating server computers when said server aggregation location data field is incomplete and said server aggregation contact strategy data field is complete.

35 9. The computer readable memory of claim 6 wherein said first instruction set includes instructions to designate said aggregation of server computers as federated server computers when said server aggregation location data field is incomplete and said server aggregation contact strategy data field is incomplete.

40 10. The computer readable memory of claim 6 further comprising a third instruction set to modify a remote procedure call parameter after said aggregation of server computers is designated.

45 11. A method of processing a remote procedure call from a client computer to an object stored on an aggregation of server computers, said method comprising the steps of:

50 checking a server aggregation location data field and a server aggregation contact strategy

data field of said remote procedure call; and designating said aggregation of server computers as replicating server computers, migrating server computers, or federated server computers based upon said checking step. 5

12. The method of claim 11 wherein said designating step includes the step of designating said aggregation of server computers as replicating server computers when said server aggregation location data field is complete and said server aggregation contact strategy data field is complete. 10

13. The method of claim 11 wherein said designating step includes the step of designating said aggregation of server computers as migrating server computers when said server aggregation location data field is incomplete and said server aggregation contact strategy data field is complete. 15

14. The method of claim 11 wherein said designating step includes the step of designating said aggregation of server computers as migrating server computers when said server aggregation location data field is incomplete and said server aggregation contact strategy data field is incomplete. 20

15. The method of claim 11 further comprising the step of processing remote procedure call information in response to said designating step. 25

16. The method of claim 15 wherein said processing step includes the step of modifying a remote procedure call parameter in response to said designating step. 30

17. The method of claim 15 wherein said processing step includes the step of compiling client/server statistics in response to said designating step. 35

18. A method executed by a client computer under the control of a program, said client computer being connected to an aggregation of server computers, said client computer including a memory storing said program, said method comprising the steps of: 40

generating at said client computer a remote procedure call to said aggregation of server computers; 45

50

55

19. The method of claim 18 wherein said designating step includes the step of designating said aggregation of server computers as replicating server computers when said server aggregation location data field is complete and said server aggregation contact strategy data field is complete.

20. The method of claim 18 wherein said designating step includes the step of designating said aggregation of server computers as migrating server computers when said server aggregation location data field is incomplete and said server aggregation contact strategy data field is complete.

21. The method of claim 18 wherein said designating step includes the step of designating said aggregation of server computers as migrating server computers when said server aggregation location data field is incomplete and said server aggregation contact strategy data field is incomplete.

22. The method of claim 18 further comprising the step of processing remote procedure call information in response to said designating step.

23. The method of claim 22 wherein said processing step includes the step of modifying a remote procedure call parameter in response to said designating step.

24. The method of claim 22 wherein said processing step includes the step of compiling client/server statistics in response to said designating step.




Figure 1
(Prior Art)




Figure 2
(Prior Art)




Figure 3




Figure 4




Figure 5





Figure 6

Aggregation Object Reference Data Structure (AORDS)**Location Information For Smart Proxy Server****Object Key****Server Aggregation Location****Location Information for Primary Server****Location Information for Secondary Servers****Server Aggregation Contact Strategy****Client Key****Figure 7**




Figure 8




Figure 9




Figure 10

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets

(11)



EP 0 817 445 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
28.03.2001 Bulletin 2001/13

(51) Int Cl.7: G06F 9/46, H04L 29/06

(43) Date of publication A2:
07.01.1998 Bulletin 1998/02

(21) Application number: 97304781.4

(22) Date of filing: 01.07.1997

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE

(30) Priority: 02.07.1996 US 678408

(71) Applicant: SUN MICROSYSTEMS, INC.
Mountain View, CA 94043 (US)

(72) Inventor: Katiyar, Dinesh
Mountain View, California 94040 (US)

(74) Representative:
Cross, Rupert Edward Blount et al
BOULT WADE TENNANT,
Verulam Gardens
70 Gray's Inn Road
London WC1X 8BT (GB)

(54) Apparatus and method for indentifying server computer aggregation topologies

(57) A method of processing a remote procedure call from a client computer to an object stored on an aggregation of server computers includes the step of checking a server aggregation location data field and a server aggregation contact strategy data field of the remote procedure call. The aggregation of server comput-

ers is designated as replicating server computers, migrating server computers, or federated server computers based upon the checking operation. Once a server computer aggregation topology is identified, parameters associated with the remote procedure call may be modified to alter the interaction with the server computer aggregation.

EP 0 817 445 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

DOCUMENTS CONSIDERED TO BE RELEVANT					
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)		
A	<p>SANTOSH K SHRIVASTAVA ET AL: "STRUCTURING FAULT-TOLERANT OBJECT SYSTEMS FOR MODULARITY IN A DISTRIBUTED ENVIRONMENT" IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, US, IEEE INC, NEW YORK, vol. 5, no. 4, 1 April 1994 (1994-04-01), pages 421-432, XP000439637 ISSN: 1045-9219 * abstract * * Chapter III. System Structure *</p> <p>---</p>	1,6,11, 18	G06F9/46 H04L29/06		
A	<p>EP 0 695 992 A (SUN MICROSYSTEMS INC) 7 February 1996 (1996-02-07) * page 9, line 31 - line 48 * * page 12, line 56 - page 16, line 3 *</p> <p>---</p>	1,6,11, 18			
A	<p>US 5 483 652 A (SUDAMA RAM ET AL) 9 January 1996 (1996-01-09) * abstract * * column 1, line 9 - line 52 * * column 5, line 39 - column 7, line 44 * * column 8, line 41 - line 46 * * column 10, line 4 - line 51 *</p> <p>---</p>	1,6,11, 18	<p>TECHNICAL FIELDS SEARCHED (Int.Cl 6)</p> <p>G06F H04L</p>		
<p>The present search report has been drawn up for all claims</p>					
Place of search	Date of completion of the search	Examiner			
THE HAGUE	7 February 2001	Kesting, V			
CATEGORY OF CITED DOCUMENTS					
X : particularly relevant; if taken alone	T : theory or principle underlying the invention				
Y : particularly relevant if combined with another document of the same category	E : earlier patent document, but published on, or after the filing date				
A : technological background	D : document cited in the application				
O : non-written disclosure	L : document cited for other reasons				
P : intermediate document	& : member of the same patent family, corresponding document				

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 97 30 4781

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

07-02-2001

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
EP 0695992 A	07-02-1996	US	5566302 A	15-10-1996
US 5483652 A	09-01-1996	JP 8115223 A		07-05-1996
NONE				

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (up to)